

**AMENDMENTS TO THE SPECIFICATION**

Please replace paragraph [0001] with the following amended paragraph:

**[0001]** While replicating only portions of the data to secondary nodes is desirable, most replication facilities are designed to copy the contents of storage locations, without regard to the type or meaning of the data contained in the storage locations. To perform an operation that recognizes the type or meaning of the data, typically application-specific software is used. For example, copying only individual files requires knowledge of the storage locations **[[are]]** included in each file, which is information that is not typically available to a replication facility. Copying an individual file is possible using a file copying utility such as xcopy, but these utilities typically do not operate on selected portions of a file. For example, if only one bit has changed in a file containing one gigabyte of data, then a file copy utility must copy the entire gigabyte of data to capture the change, which is also very time consuming. A faster way to restore and/or synchronize selected data from large volumes of data and/or files is needed.

Please replace paragraph [0038] with the following amended paragraph:

**[0038]** While application 115A and storage manager / replicator 120A may run on the same computer system, such as primary node 110A, the hardware and software configuration represented by primary node 110A may vary. Application 115A and storage manager / replicator 120A may execute on different computer systems. Furthermore, storage manager / replicator 120A can be implemented as a separate storage management module and a replication module that operate in conjunction with one another. Application 115A itself may **[[have]]** provide some storage management functionality.

Please replace paragraph [0045] with the following amended paragraph:

**[0045]** The set of one or more storage locations can be represented as **a** set of one or more extents. A file extent includes a layout of physical storage locations on a physical storage volume. The file extent typically includes an address for a starting location in the file and a size (the number of contiguous locations beginning at the address). A single

file can include several non-contiguous portions (each of which will have a respective starting location and size). One of skill in the art will recognize that file extents can be expressed in storage units such as file clusters, but are referred to herein as locations on the volumes for simplicity purposes.

Please replace paragraph [0049] with the following amended paragraph:

[0049] Fig. 3B shows an example of data replicated using volume sieves. Sieve 320A includes a property having an operation of replication to replication volume #1 (replication volume 310A), which applies to the set of locations beginning at location 7 and including three locations. In this example, sieve 320A applies to storage locations 7, 8, and 9 of region R3, having respective values 'q', 'C,' and '@.'